

*NASA-CR-190517*  
*MIT-18-CR*

**SPACECRAFT DYNAMICS  
CHARACTERIZATION AND CONTROL  
SYSTEM FAULT DETECTION**

**Final Report**

**Volume 1**

**SUMMARY AND CONCLUSIONS**


*103021*  
*P-54*

(NASA-CR-190517) SPACECRAFT DYNAMICS  
CHARACTERIZATION AND CONTROL SYSTEM FAILURE  
DETECTION, VOLUME 1 Final Report 9 Feb. 1989  
- 8 Feb. 1992 (MIT) 54 p


N92-29422

Unclass

G3/18 0109021



**DEPARTMENT OF  
AERONAUTICS AND ASTRONAUTICS  
MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
CAMBRIDGE, MASSACHUSETTS 02139**



1

ND-61-705  
IN-13-OR

**SPACECRAFT DYNAMICS  
CHARACTERIZATION AND CONTROL  
SYSTEM FAULT DETECTION**

**Final Report**

**Volume 1**

**SUMMARY AND CONCLUSIONS**

109021

P-54

# **SPACECRAFT DYNAMICS CHARACTERIZATION AND CONTROL SYSTEM FAILURE DETECTION**

Final Report

NASA Research Grant No. NAG-1-968

**Volume I**

## **SUMMARY AND CONCLUSIONS**

Mark S. Webster  
Christiaan M. van Schalkwyk  
Wallace E. Vander Velde

*Department of Aeronautics and Astronautics  
Massachusetts Institute of Technology  
Cambridge, Massachusetts 02139*

February 10, 1992

## SPACECRAFT DYNAMICS CHARACTERIZATION AND CONTROL SYSTEM FAILURE DETECTION

### Introduction

This is the final report of work accomplished under NASA Research Grant No. NAG-1-968 during the period February 9, 1989, to February 8, 1992. The work under this grant has been directed to two aspects of the control of flexible spacecraft:

- The modeling of deployed or erected structures including nonlinear joint characteristics
- The detection and isolation of failures of the components of control systems for large space structures

The motivation for the first of these research tasks is the fact that very large assemblies in space will have to be built or deployed in situ. A likely scenario is, in fact, a combination of these wherein modules which are folded for transportation into orbit are erected to their final configuration and then joined with other such erected modules to form the full assembly. Any such erectable modules will have joints - possibly many of them. Joints designed for existing experimental structures, such as the Mini-Mast at NASA Langley Research Center, display distinctly nonlinear load-displacement behavior. It remains to be seen whether or not joints designed for operational assemblies will have nonlinear properties, but it seems prudent to develop a methodology for dealing with that possibility.

The motivation for the second of these research tasks is the fact that we foresee large assemblies in space which will require active control to damp vibrations and/or hold a desired shape. Lightweight structures will be very flexible, with many elastic modes having very low frequencies. In order to control these modes well, the control system will likely require many sensors and many actuators, probably distributed over much of the structure. Also, it would be prohibitively expensive to visit these assemblies often for the purpose of repairs and servicing. The combination of a large number of control system components with long operational periods virtually guarantees that these systems will suffer control system component failures during operation. The control system must be designed to tolerate failures of some sensors and actuators - and be able to continue to perform its function. A central feature of most fault tolerant control system architectures is detection of component failures and

isolation of the fault to a single component. The failure detection and isolation (FDI) function is therefore central to highly reliable control systems and is one of the fundamental technologies which must be demonstrated on ground-based experimental facilities and then in space flights.

The work accomplished on each of these research tasks will be summarized briefly in this volume. Volumes II and III are more detailed presentations of the theory and results in each of these areas.

### **Modeling Beam-Like Space Trusses With Nonlinear Joints**

**Approach** - Large space structures are normally characterized by high order finite element models. These models have a large number of degrees of freedom and usually must be abbreviated in some way for the purpose of identification and control system design. The most common approaches to model reduction select a subset of the modes produced by the finite element model; different methodologies employ different criteria for selection of the important modes. A different approach recognizes that the global behavior of a truss should resemble that of a continuum beam with properly chosen beam properties. The appropriate beam stiffness can be found by equating the strain energy in the finite element representation of the structure and the approximating beam for corresponding displacements.

A still better representation of a truss results from approximating each bay of the truss with a beam element. In the context of this work, this modeling approach has an additional advantage in providing a more direct indication of the displacements of the nodes of the truss which are required to compute the forces across nonlinear joints. Therefore, except for some initial comparisons of accuracy with linear joints, all of the results in this research were generated with the truss dynamics represented by a beam element for each cell of the truss.

This approximation of complex trusses by beam elements achieves the objective of reducing the number of degrees of freedom in the model of the truss dynamics. An additional methodology is required to account for nonlinear behavior of the truss. In this work we have concentrated on nonlinear joint properties for the reasons cited in the Introduction. Nonlinear structural behavior is most readily recognized in sinusoidal response tests. Resonant frequencies that vary with forcing amplitude, and sinusoidal response characteristics that differ for positive-going sine sweeps and negative-going sine sweeps are classical indicators of nonlinear behavior. An approximate representation of the

dominant behavior of a nonlinear element under sinusoidal excitation is given by the sinusoidal-input describing function. The advantage of employing describing function characterization of the nonlinear elements rather than the nonlinear functions themselves is that this permits analytic (actually computational) solution for properties such as sinusoidal response as opposed to being limited to direct simulation. The describing function for a nonlinear element represents exactly the fundamental component of the output of the element under sinusoidal excitation. The approximation made in using describing functions is that the inputs to all the nonlinear elements are assumed to be sinusoids whereas in fact some of the higher harmonics generated by the nonlinearities propagate to the inputs of the nonlinear elements. The accuracy of the method, then, depends on the nature of the nonlinear characteristics in terms of how much harmonic distortion they generate, and on the degree of attenuation of the higher harmonics through the linear dynamics of the system.

The approach taken in this research to modeling beam-like trusses with nonlinear joints is to combine describing function representation of the joint characteristics with beam element representation of the truss dynamics for the purpose of reducing the number of degrees of freedom in the model.

**Results** - The accuracy of various reduced dof models of a truss is compared first for the case of a linear truss. The example is a planar cantilevered truss with a varying number of bays. The computed frequencies of the first transverse bending mode are plotted in Figure 1. The full finite element model of the truss is considered the reference and the approximate models are compared with that. The most obvious trend is that the accuracy of all the models improves with number of bays. This simply reflects the fact that the approximations are based on beam properties and the truss resembles a beam more as its length-to-width ratio increases. For shorter trusses (less beam-like trusses) the beam element models are much more accurate than the full continuum beam models - whether Bernoulli-Euler theory or Timoshenko theory is applied to the approximating beam.

To display better the accuracy of the two beam element models, the percent error in predicting the first mode frequency is shown for these two models in Figure 2. The two models differ in the number of dof which are represented for each beam element. The four dof model represents transverse displacement and rotation at each end of the element. The six dof model has in addition an axial dof at each end. The additional dof is clearly helpful with the error in first mode frequency for that model being less than 5 percent for all truss lengths. But the four dof model




Figure 1. First Frequency for a Cantilevered Truss with Varying Number of Bays




Figure 2. Percentage Error of Equivalent Beam Models From Truss Finite Element Model



also predicts the first mode frequency within 5 percent if the truss has more than 6 bays.

To account for nonlinear joint characteristics, each strut of the truss is represented as shown in Figure 3. The strut is shown attached to a joint at each end. The joint can have linear stiffness and damping characteristics in addition to possible nonlinear properties. The elasticity of the strut is also accounted for and the force-displacement property of this whole assembly is represented by a describing function. This gives strut stiffness and damping which vary with the amplitude of sinusoidal displacement across the joint-strut-joint element.

With the truss struts modeled in this way, the whole truss model has the same structure as that of a finite element model, but elements of the stiffness and damping matrices are functions of the amplitudes of the displacements across various struts. Those displacements, in turn, depend on the stiffness and damping properties. So a numerical solution is required to calculate sinusoidal response properties of a structure. The conditions for harmonic balance under steady state response to sinusoidal forcing can be written as a set of equations to be satisfied. A Newton-Raphson iteration was employed to solve these equations using finite difference approximations to the required derivatives.

To illustrate the results one can obtain with this modeling, the planar 5 bay truss pictured in Figure 4 has been studied with several types of nonlinear joint properties. The degrees of freedom for which responses are shown in the following figures are identified in this figure. As a reference for comparison with the nonlinear cases, the responses of the 6 identified dofs to transverse forcing at the end of the truss are shown in Figures 5 and 6 for the case of linear joints. The responses to 4 different levels of forcing are shown. As must be true of a linear model, the resonant frequencies and damping characteristics are independent of the level of response.

The response of the truss with hardening spring joint characteristics is shown in Figures 7 and 8. In the linear range, the spring stiffness is the same as that of the linear model whose response is given in Figures 5 and 6. The stiffness in the upper range of deflections is twice that of the lower. Also shown in Figures 7 and 8 along with the responses of the reference dofs are the so-called backbone curves. These curves are found from a simplified form of the dynamic equations and are intended to indicate the resonances. These curves are often the clearest indicators of nonlinear behavior. It can be seen that at low amplitude, these backbone curves indicate the same




Figure 3. Joint-Strut-Joint Schematic


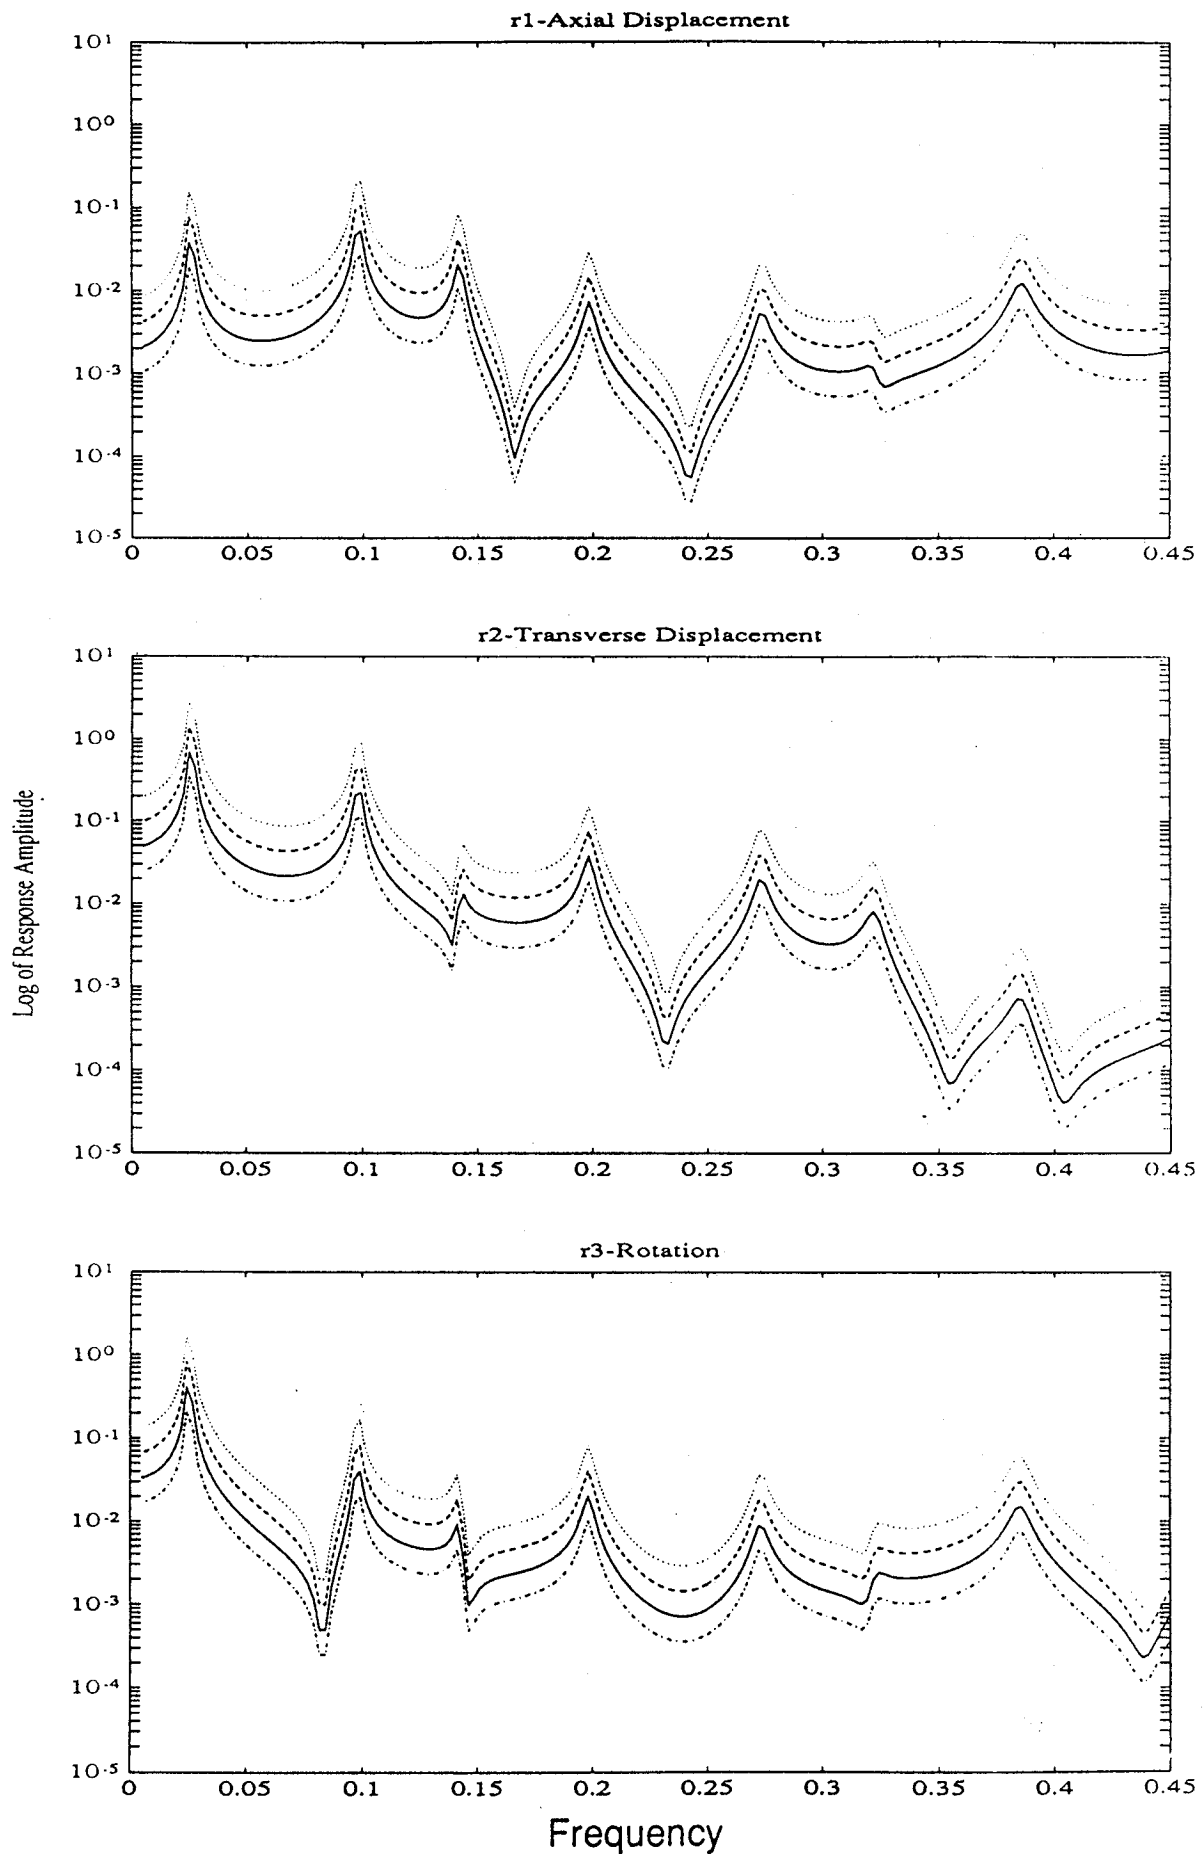
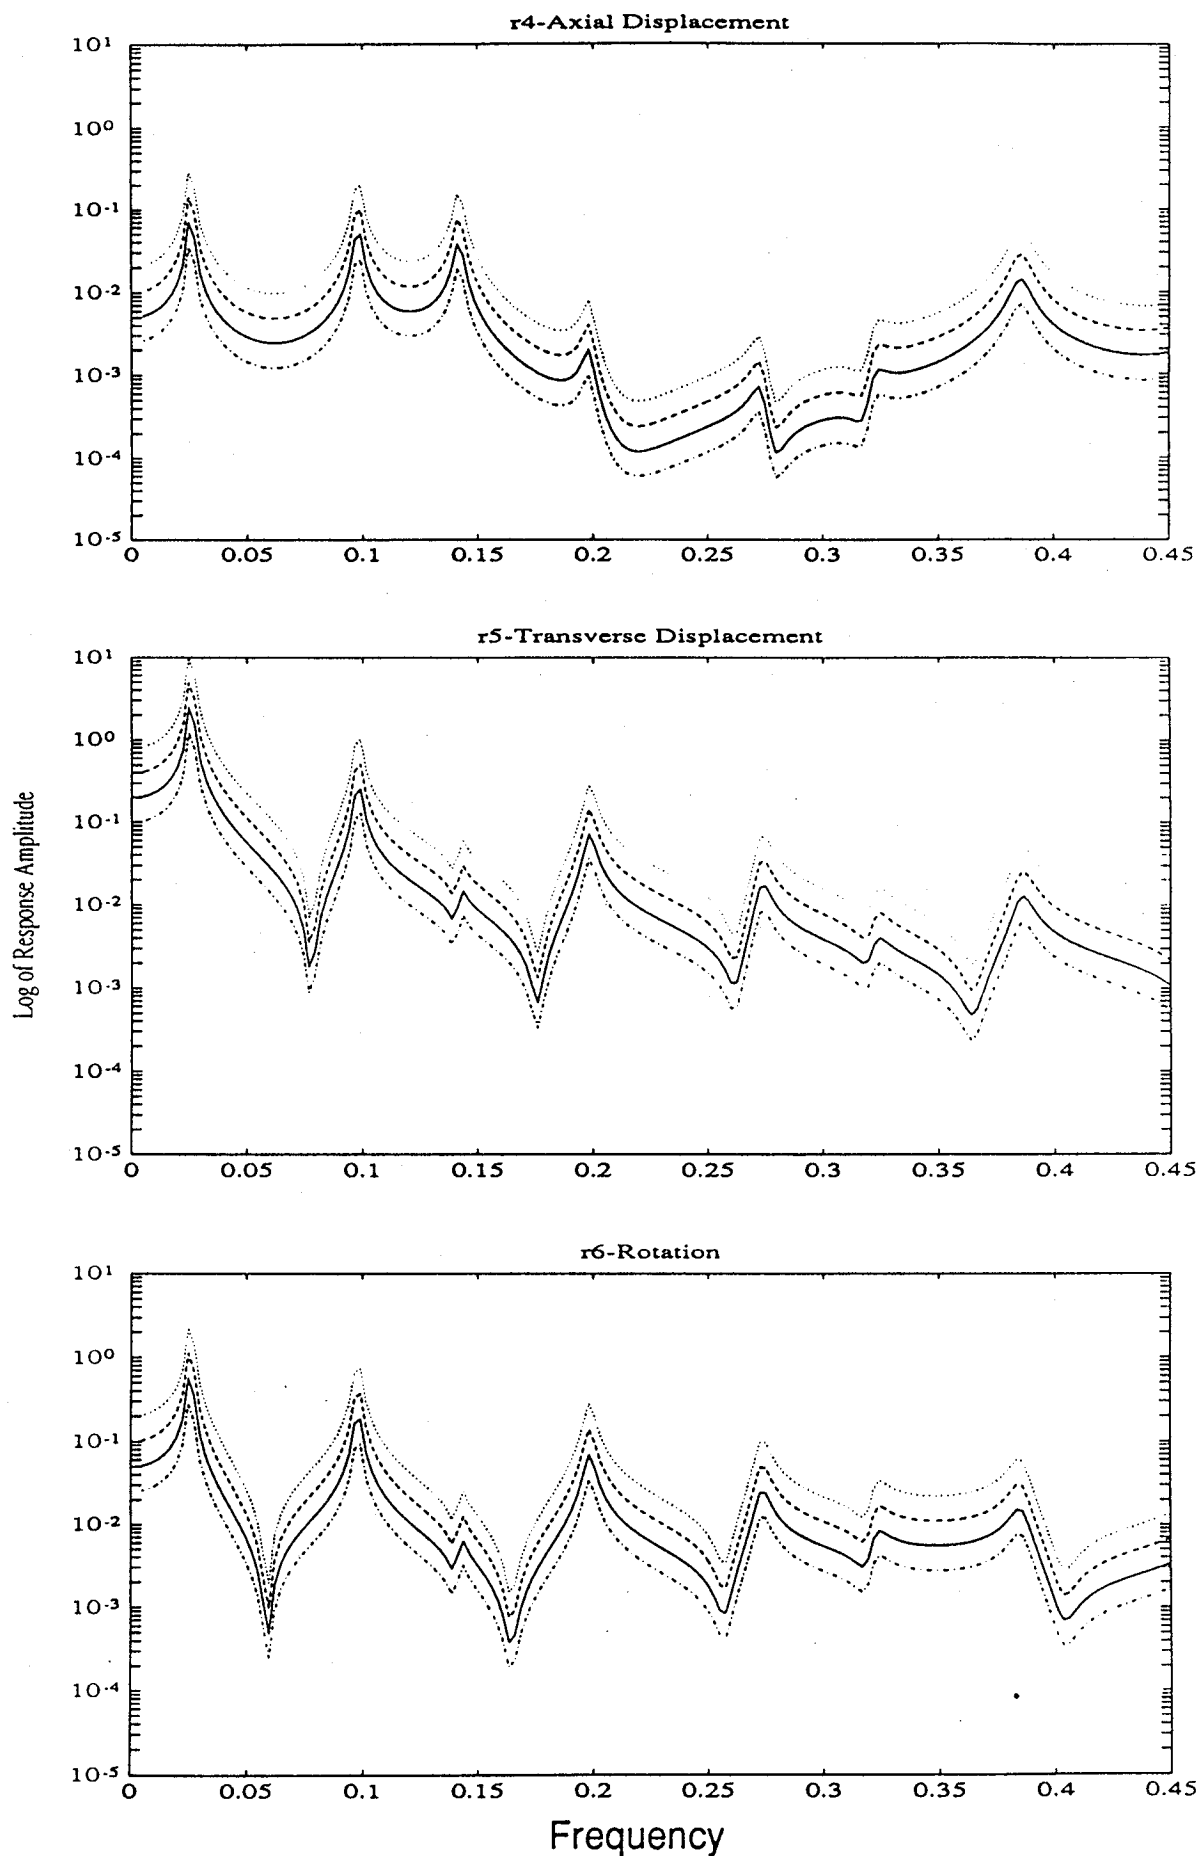


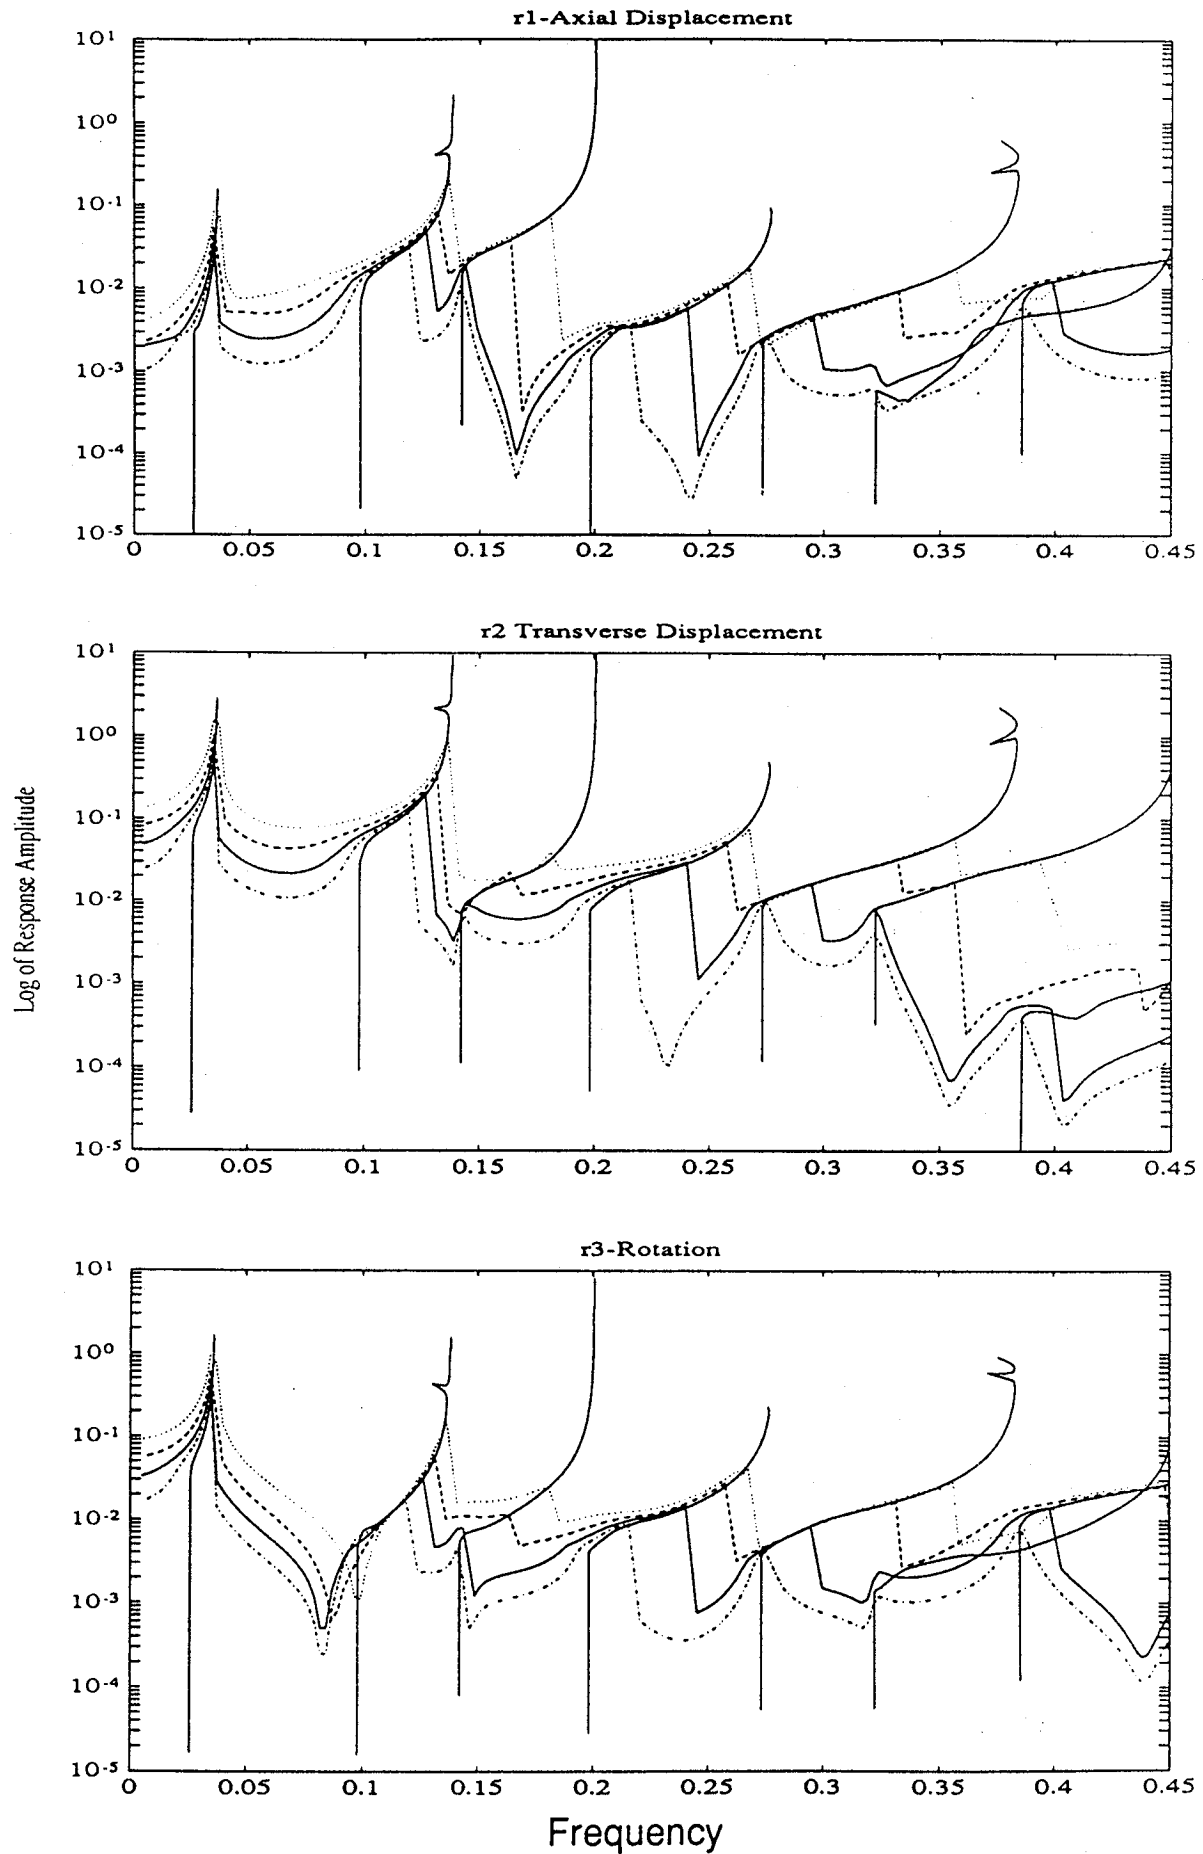
Figure 4. 2-D Five Bay Cantilevered Truss with Equivalent Beam Model



**Figure 5.** Response at End of Bay 2 to Transverse Sinusoidal Forcing at tip of Five Bay Truss with Linear Joints Modeled by Equivalent Finite Element Beam



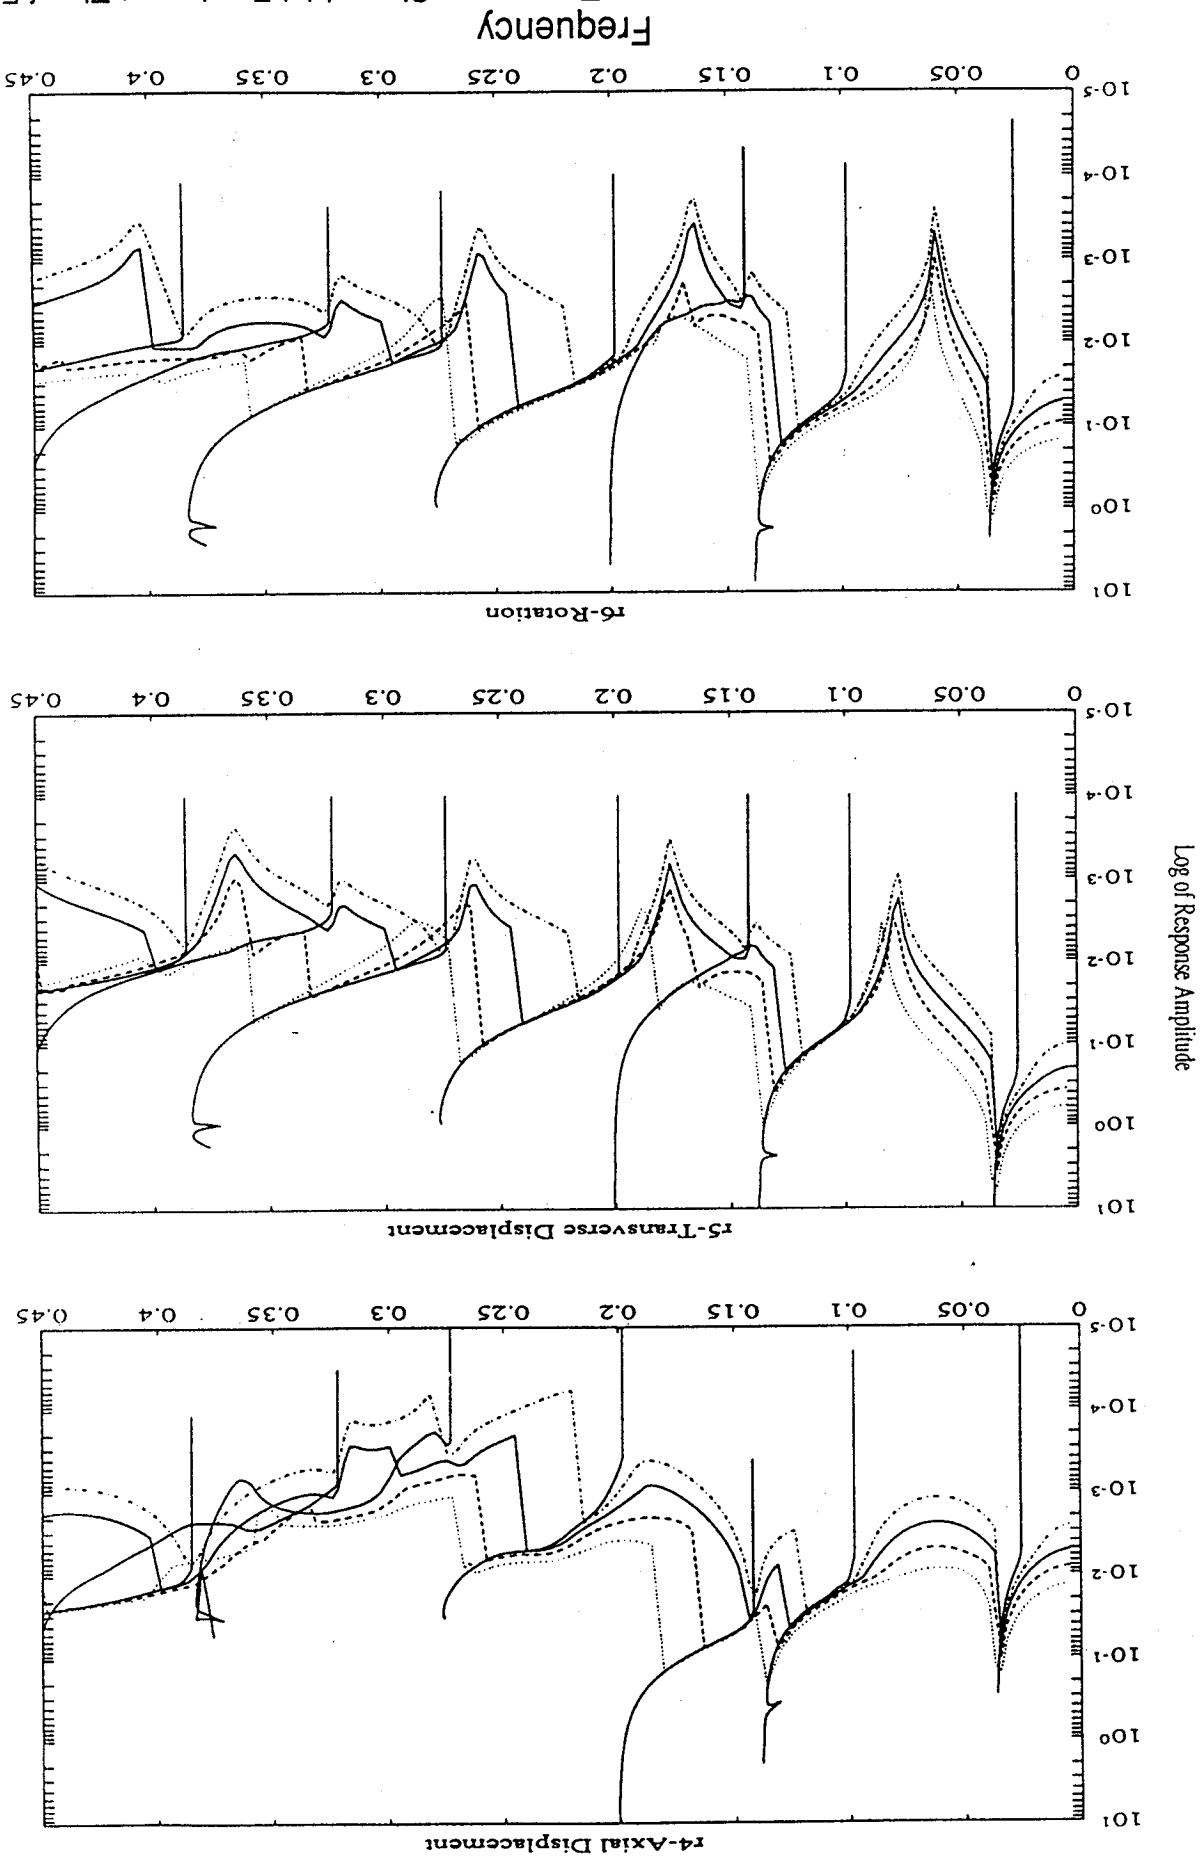
**Figure 6.** Response at End of Truss to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Linear Joints Modeled by Equivalent Finite Element Beam



**Figure 7.** Response at End of Bay 2 to Transverse Sinusoidal Forcing at tip of Five Bay Truss with Hardening Gain Change Joints Modeled with Nonlinear Equivalent Beam

Figure 8. Response at End of Truss to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Hardening Gain Change Joints Modeled by Nonlinear Equivalent Beam

10

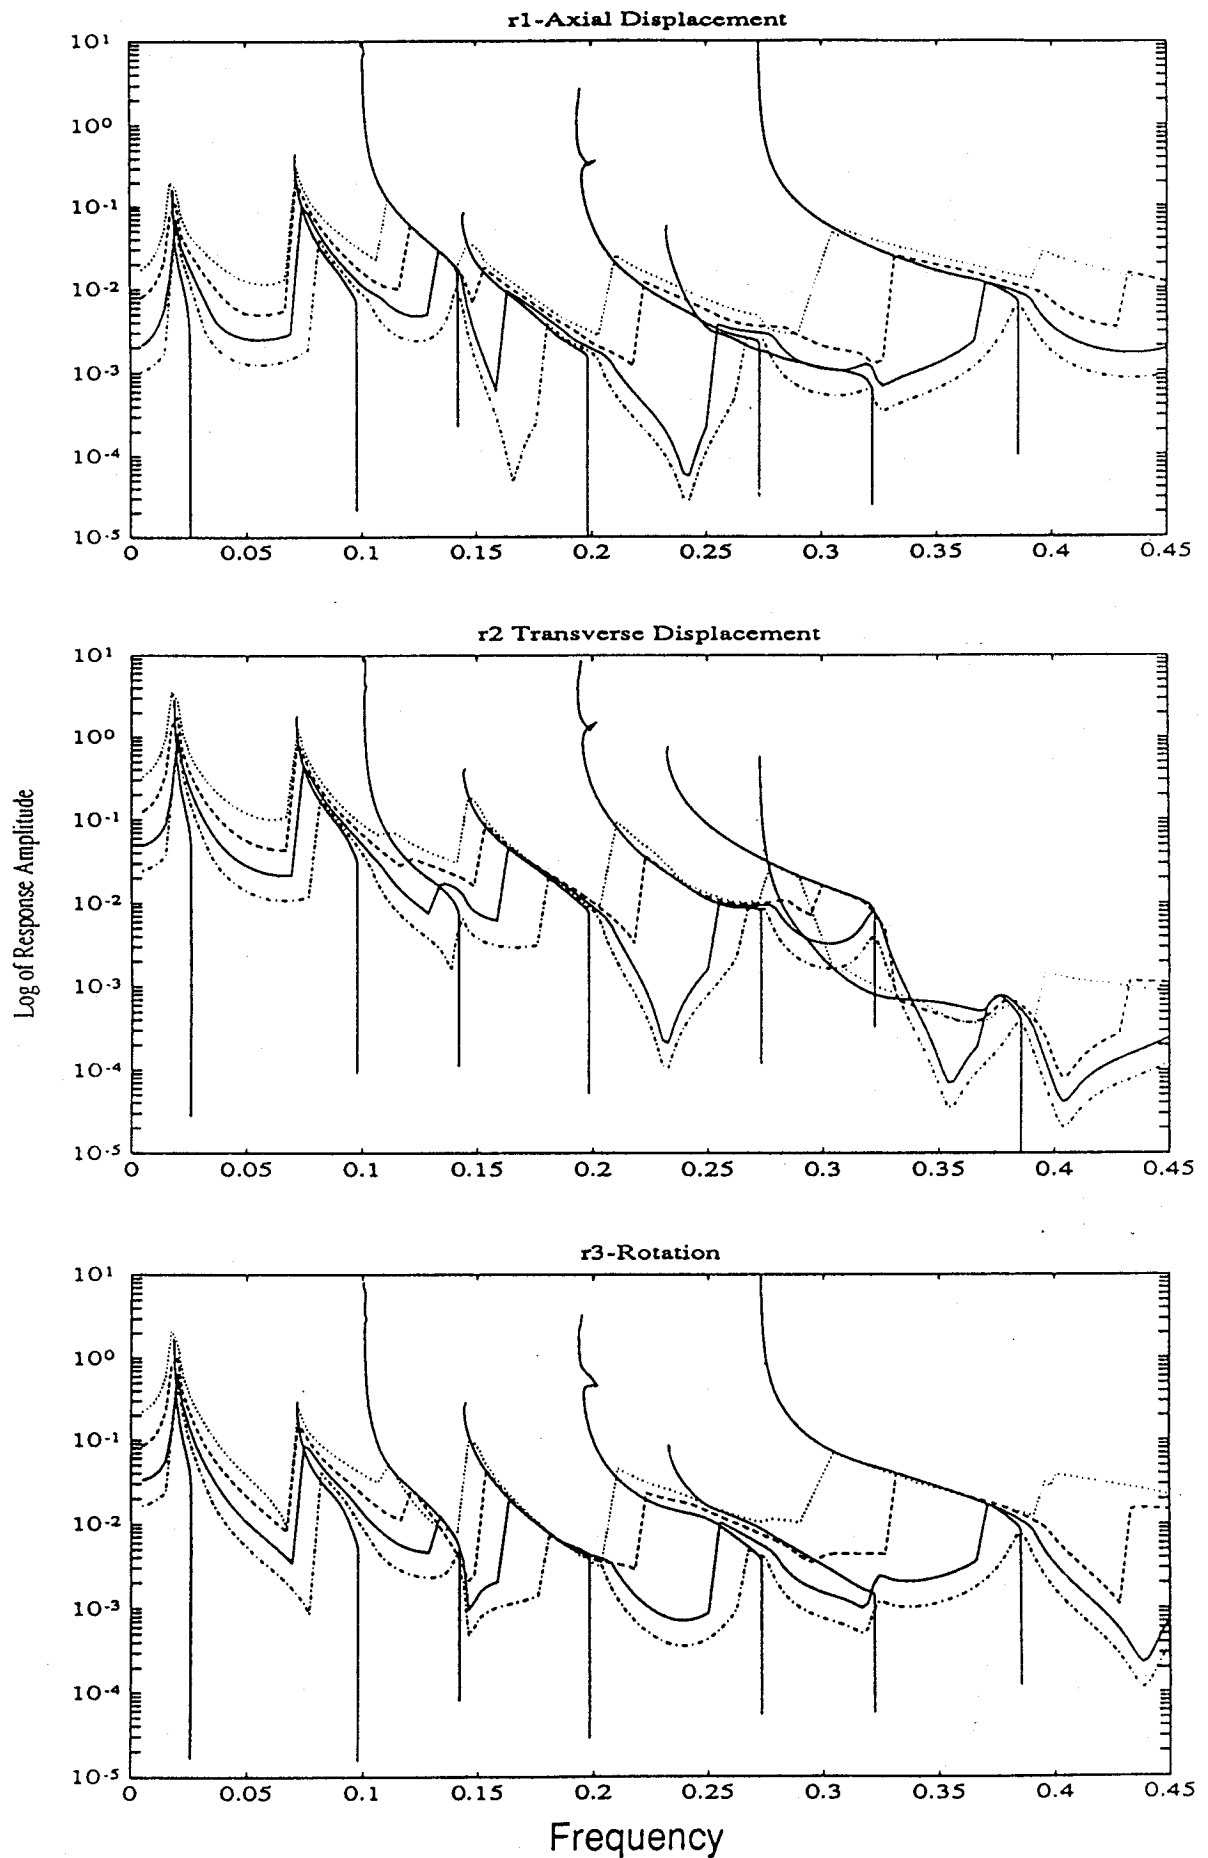


resonant frequencies as those of the linear model. At higher amplitudes, as the joints begin to experience nonlinear behavior, the resonant frequencies move higher. Asymptotically, the resonant frequencies shift higher by a factor of the square root of 2, consistent with the spring stiffness ratio of 2. At high amplitude, some of the backbone curves show small "glitches" which are due to numerical problems and do not indicate predicted behavior of the truss. In some areas the response curves jump discontinuously from one response curve to another rather than tracing out a smooth profile as the forcing frequency is changed. This is indicative of a region of multivalued response and the phenomenon is known in the nonlinear mechanics literature as "jump resonance". If one were to retrace the frequency response curve in the opposite direction, one would likely see the other branch of the response curve over some interval.

Figures 9 and 10 show the responses of the truss with softening spring joints. Again, the joint stiffness is the same as that of the linear joints in the linear range. At large deflections, the stiffness is one half that at low deflections. These response curves are very similar to those with hardening spring joints except that the resonant frequencies move lower at higher amplitude rather than higher. Again, there are clear indications of jump resonance.

As a final example of the response of the five bay truss, the joints are considered to be pinned joints with some free play between the pins and the holes. There is also friction between the two end pieces of the joint. An idealization of the resulting force-displacement relation is shown in Figure 11. The joint has a linear stiffness characteristic everywhere except in an intermediate region around the origin. When the force across the joint reaches the friction level, the pin slips to the opposite end of the region of free play. Then the joint experiences its original stiffness again. As the force is reversed in a sinusoidal cycle, the force must reach the friction level in the opposite direction before the pin slips again. The stiffness of the joint is the same as that of the linear joint. The responses of the reference dofs of the truss are shown in Figures 12 and 13. In this case, the resonant frequencies are the same as for the linear truss at both low and high amplitudes when the joints remain in a linear range or when they are beyond the slipping region over most of the cycle. At intermediate amplitudes the resonant frequencies are lower as the slipping of the joints reduces the effective stiffness. This joint characteristic also induces damping, and the resonant peaks are seen to be lower and broader than in the previous examples.


A particularly difficult aspect of dealing with




**Figure 9.** Response at End of Bay 2 to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Softening Gain Change Joints Modeled by Nonlinear Equivalent Beam



Figure 10. Response at End of Truss to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Softening Gain Change Joints Modeled by Nonlinear Equivalent Beam





**Figure 11. Sliding Pin Nonlinearity**

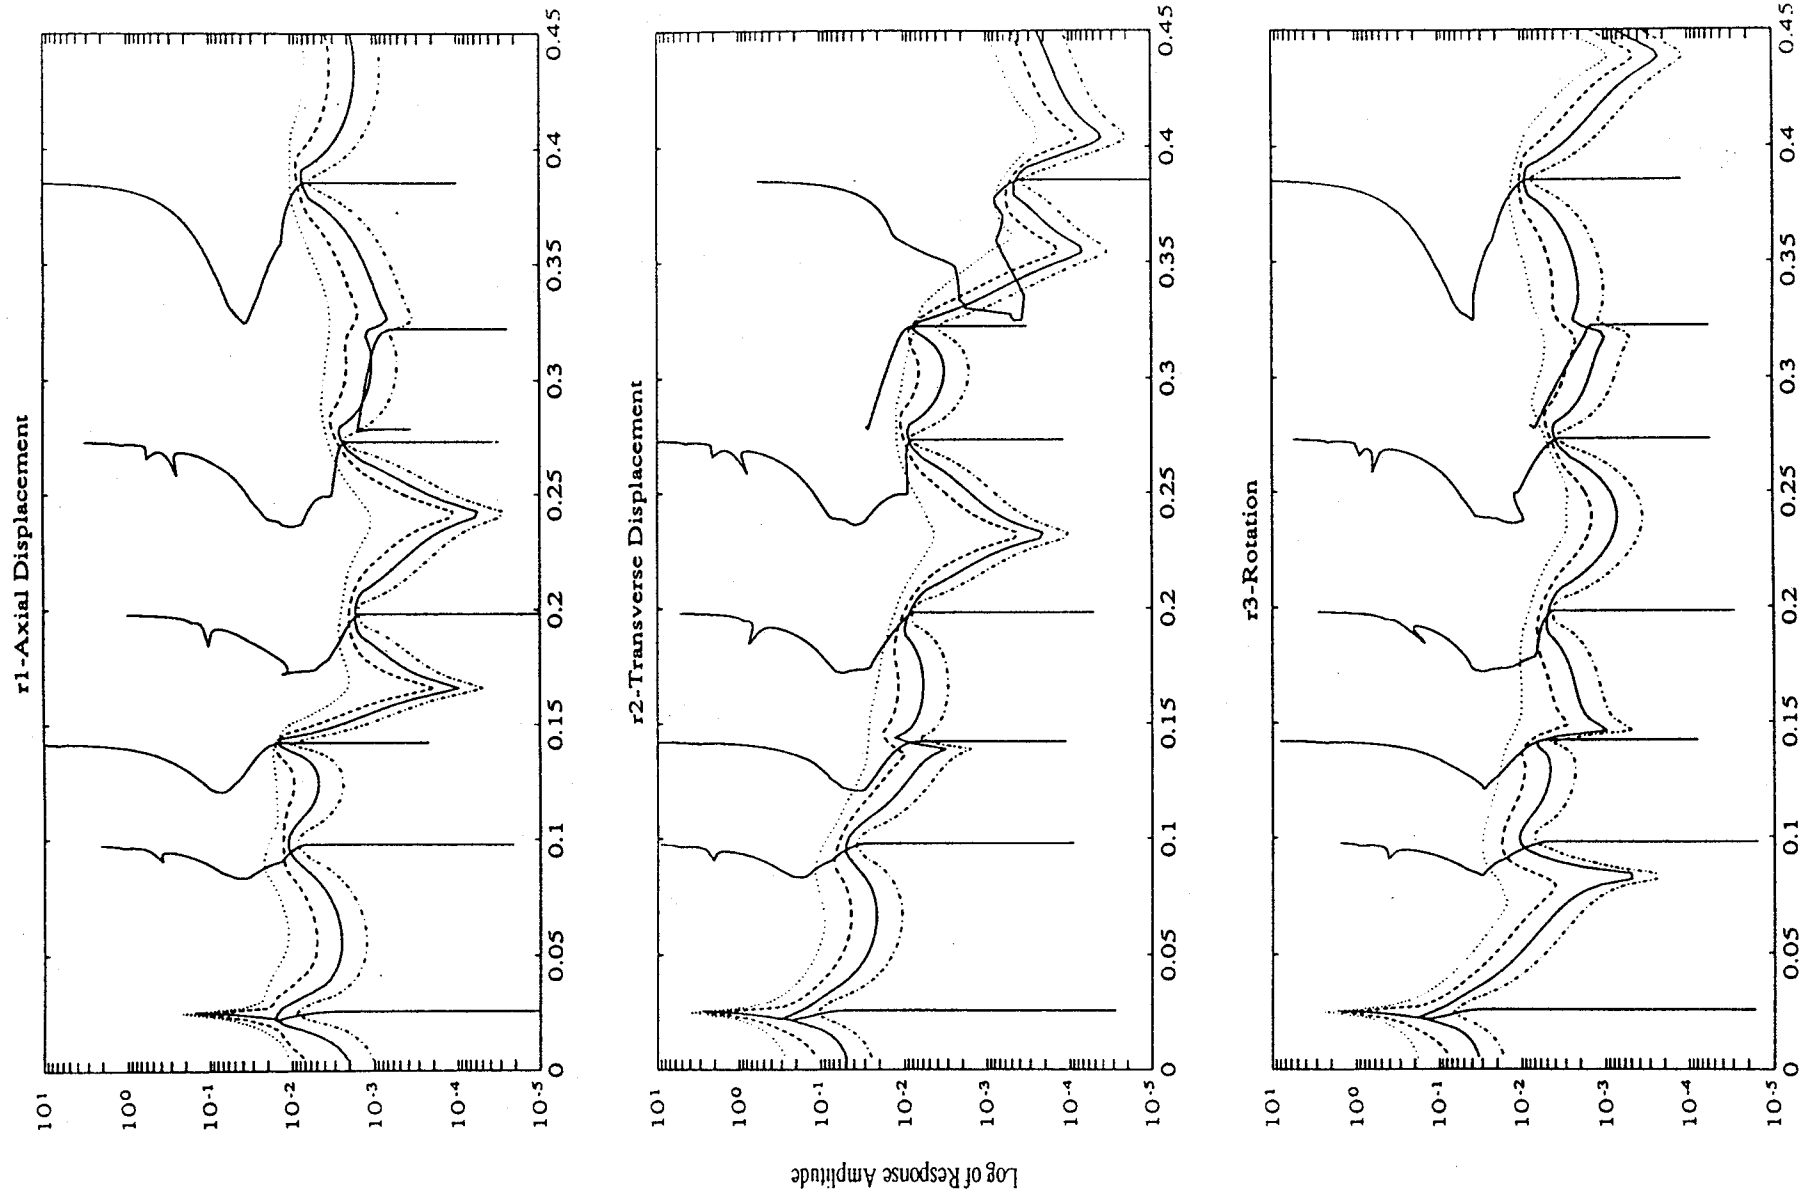
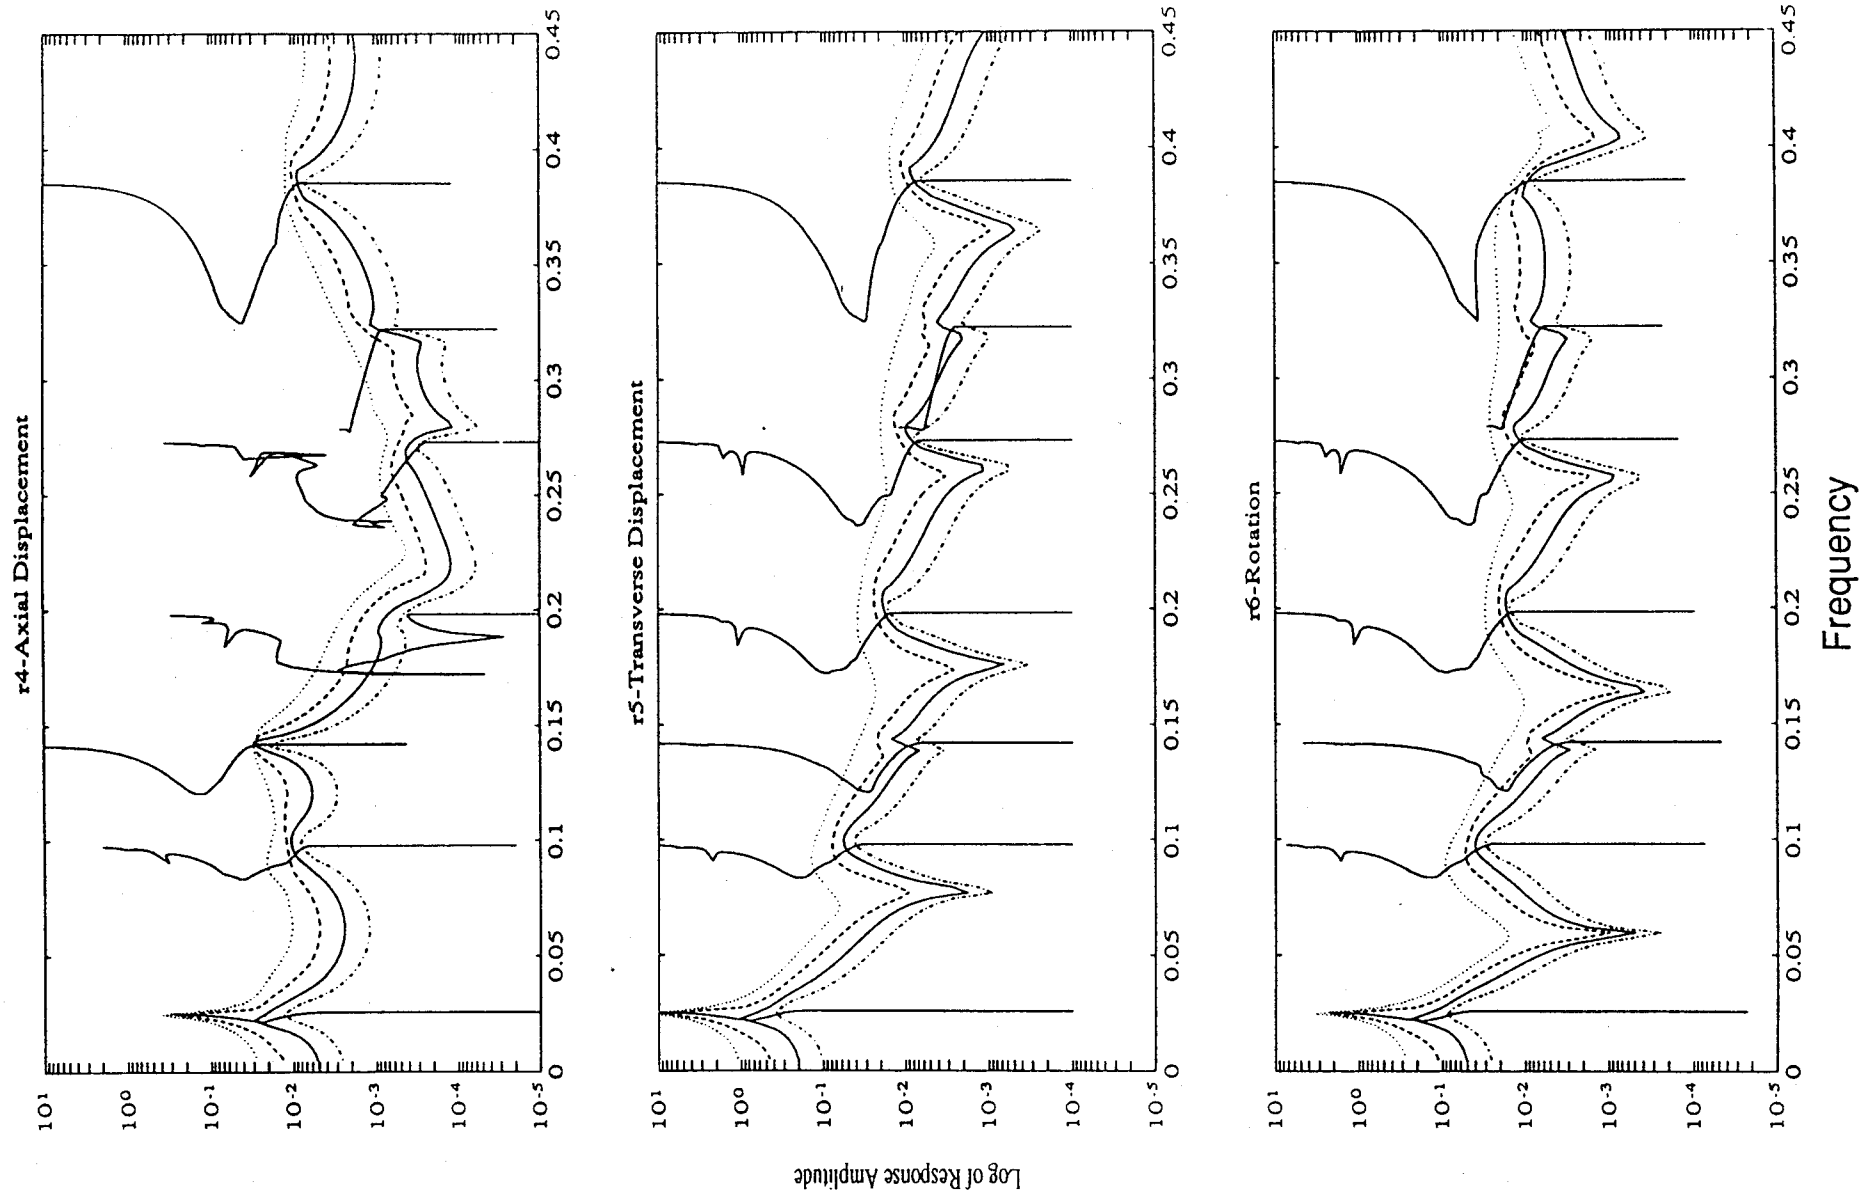


Figure 12. Response at End of Bay 2 to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Sliding Pin Joints Modeled by Nonlinear Equivalent Beam




**Figure 13.** Response at End of Truss to Transverse Sinusoidal Forcing at Tip of Five Bay Truss with Sliding Pin Joints Modeled by Nonlinear Equivalent Beam


nonlinear structures is the fact that their dynamic behavior changes with any preload across the joints as might be induced by the weight of the structure if tested on the ground, or by gravity gradient or spinning the structure. An extension to the usual sinusoidal-input describing function theory can account for such effects. This is the dual-input describing function wherein one accounts for a bias plus a sinusoid at the input to the nonlinearity. Then the describing function models the transmission of the bias and the sinusoid through the nonlinear element. The gains to both the bias and the sinusoid depend on the amplitudes of both the bias and the sinusoid.

As an illustration of the effect of a preload on a structure with nonlinear joints, Figures 14 and 15 show backbone curves and response curves for one of the reference dofs of the five bay structure with hardening spring joints with and without a preload. The no preload case is the same as in Figure 8. The preload has been simulated as being induced by operating the truss in a vertical orientation on the ground. The load is thus due to the weight of the truss itself. A relatively heavy truss has been considered such that the weight has put some of the joints in the nonlinear region. Therefore the backbone curves show a shift in the resonant frequencies even for very small sinusoidal forcing. As the amplitude of the response to the sinusoidal forcing becomes large enough, the two cases converge because the joints are in their large amplitude condition most of the time regardless of the preload. The marked differences between the two sinusoidal response curves emphasize the importance of accounting for any preload existing on a structure if it has any nonlinear elements.

The final example deals with a real structure - the Mini-Mast test article at the NASA Langley Research Center. This mast is a deployable space truss erected in a vertical orientation. Being deployable, most of the struts which comprise the truss are mounted to pinned joints at both ends. The truss can be driven by three torque wheel actuators mounted on a platform at the top, and it is heavily instrumented along its length. Figure 16 shows the Mini-Mast, which is 20 meters high, installed in a high bay facility at Langley. Figure 17 shows the mast in a partly deployed state. The fact that the Mini-Mast displays nonlinear behavior is clear from sinusoidal response testing as shown in Figures 18 and 19. The nonlinearities, which are believed to be mostly in the joints, obviously act in a complex way as the x axis response (lateral bending) displays a softening spring characteristic with decreased damping as the amplitude increases whereas the z axis response (torsion) displays a softening spring characteristic with increased damping as the amplitude increases.



**Figure 14.** Comparison of Backbone Curves of Five Bay Truss with Hardening Gain Change Joints with Large Pre-load (solid line) and No Pre-load (dotted line)



**Figure 15.** Comparison of Sinusoidal Response Curves of Five Bay Truss with Hardening Gain Change Joints with Large Pre-load (solid line) and No Pre-load (dotted line)




Figure 16. Mini-Mast Test Facility at NASA/Langley Research Center






Figure 17. Partially Deployed Mini-Mast





**Figure 18.** Transfer Function from X-Axis Torque Wheel to Tip Displacement Sensor in Vicinity of First Bending Mode, Mini-Mast Experimental Data



**Figure 19.** Transfer Function from Z-Axis Torque Wheel to Tip Displacement Sensor in Vicinity of First Torsion Mode, Mini-Mast Experimental Data

A beam element approximate model of the Mini-Mast was assembled in which a single beam element is used to represent each bay of the mast - of which there are 18. Using linear joint models, this reduced dof model matched the experimentally observed first mode frequencies in bending and torsion very well. Unfortunately, there was not available any data on the force-displacement properties of the joints in the mast. So the joints could not be modeled directly from experimental or other data; their properties had to be guessed. The nonlinear behavior is most conspicuous in torsion, so a nonlinear characteristic was defined which has a softening property and hysteresis to add damping. This nonlinearity was called a "natural" nonlinear joint and it is pictured in Figure 20. It does not have a direct physical motivation, but even if experimental joint properties were available, their interpretation as installed in the mast would be complicated by the strange geometry required for compact folding of the structure. The computed response of the mast to sinusoidal forcing in torsion is shown in Figure 21. These responses have the same character as the experimental results shown in Figure 19. The response curves show decreased resonant frequency and increased damping with increasing amplitude. The calculated curves, especially for the higher excitation levels, show sharper jumps than do the experimental curves. This is indicative of multi-valued response where the calculated results jump sharply from one response mode to the other. The experimental results display multiple response modes too, but the transitions from one mode to the other are smoother - as if the transition is being smoothed by noise or disturbances. The frequency at which the transition occurs is not something that one would expect to match. The jump occurs somewhat at random depending on prior conditions. In both the calculated and experimental results, for example, the transition has a definite dependence on the direction of the sinusoidal sweep.

All of the cases described above have been examples of open loop structural responses to sinusoidal excitation. The same methodology: beam element modeling of the structural dynamics and describing function modeling of the nonlinear joints, can be applied to controlled structures as well. The controllers are normally designed on the basis of a linear model of the structure, and this approach can be used to evaluate the stability and performance of the controlled structure with its nonlinear characteristics accounted for. This procedure has been illustrated for several types of controllers on a simple example structure. The controllers evaluated are a Linear Quadratic Regulator - which requires full state feedback, a robust version of an LQR controller, a Linear Quadratic Gaussian controller - which is a model based compensator not requiring full state measurements, and an H infinity controller - which is another model based compensator. The disturbance rejection




Figure 20. Natural Nonlinear Joint




Figure 21. Transfer Function from Z-Axis Torque Wheel to Tip Displacement Sensor in Vicinity of First Torsion Mode, Mini-Mast Nonlinear Equivalent Beam Analytical Model with Natural Joints

transfer functions for the system with the  $H$  infinity controller are shown in Figure 22. The controller was designed on the basis of a linear model which used the small-amplitude stiffness of the joints, and the performance is evaluated for joints having both softening and hardening spring characteristics. The nonlinear joints clearly have an appreciable effect on these transfer functions. The response peaks move up with hardening joints and down with softening joints as one would expect. The peak amplitude is also somewhat worse with the nonlinear joints - indicating degraded disturbance rejection.

## Detection and Isolation of Control System Component Failures

**Approach** - Many approaches to control system failure detection and isolation (FDI) have been reported in the literature. Some are applicable only to sensors and not to actuators. An example of this is the use of several Kalman filters to estimate the state of the system using different subsets of the sensors. If a sensor fails, only the filters which use that measurement would be affected. Other methods are designed to test for signatures of specific failures - including the mode of failure. All the hypothesis test methods, including the Generalized Likelihood Ratio method, fall into this category. The test must be designed for a particular mode of failure such as a bias shift or an increase in variance. Although the test for a particular failure mode can be made optimally sensitive, one must wonder how many failure modes to design the FDI system for.

Among all the FDI methodologies known to the authors, only two share the advantages of being applicable both to sensor and actuator monitoring, and do not depend on specification or hypothesis of modes of failure. This means that the same methodology can be used to monitor both of the major classes of components of control systems for large space structures - sensors and actuators - and any abnormal behavior of these devices will be sensed without having to guess the manner in which they will fail. These two FDI methodologies are generalized parity relations and the failure detection filter. These are the approaches that have been considered in this research.

The objective of this research task was to demonstrate the detection and isolation of failures of components for control systems for flexible spacecraft using experimental facilities already operational. These facilities are the Mini-Mast at NASA Langley Research Center and the ACES facility at the Marshall Space Flight Center. The typical structure of an FDI system is shown in Figure 23. Most of the effort in this work has been directed to the first stage of this FDI processor - the generation of a residual signal to be monitored for the signature of a failure. These

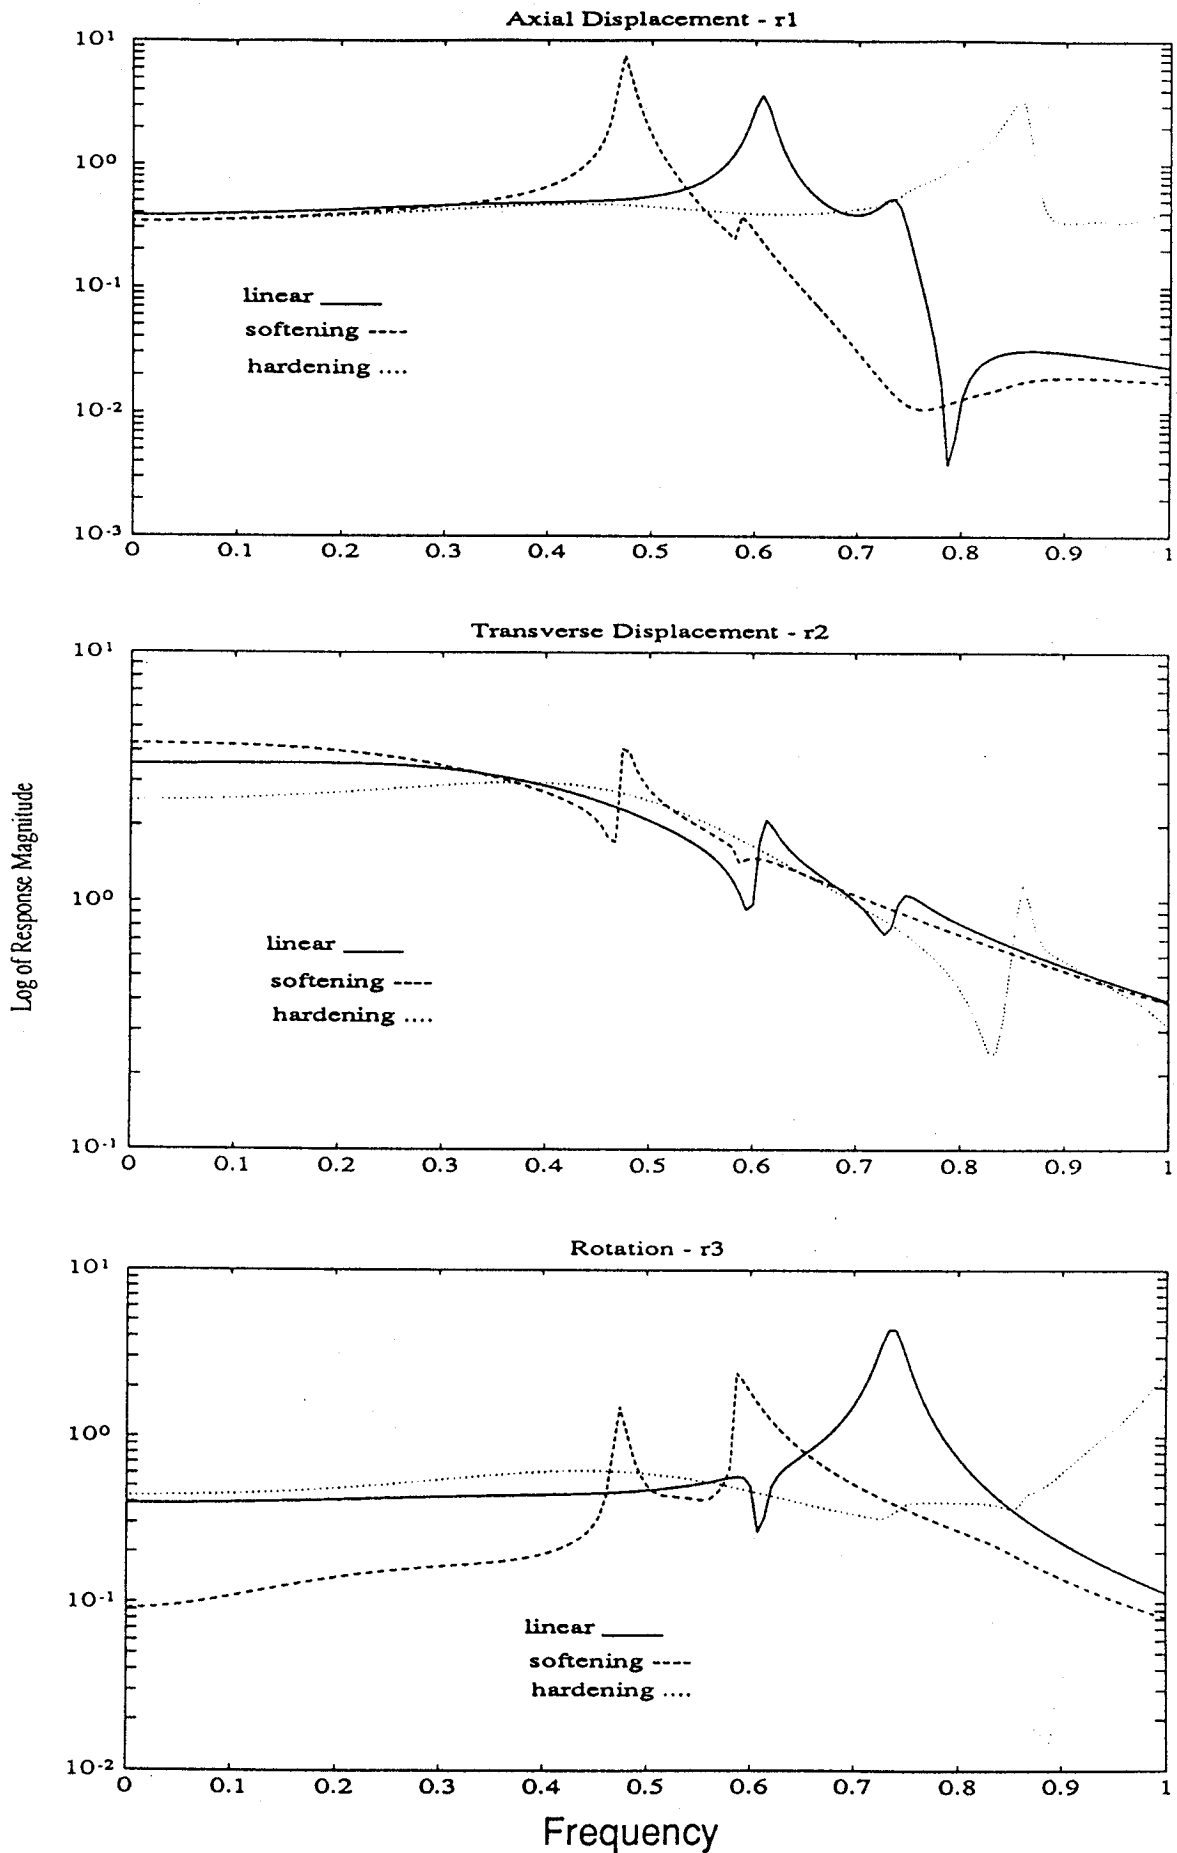


Figure 22. Transfer Function from Tip Excitation to Displacements for Structure with Gain Change Joints with H Infinity Compensator with Full Control




Figure 23. FDI block diagram.

residuals must have the property that, except for noise and other unmodeled effects, they are zero under normal system operation. When a sensor or actuator fails, one or more of the residuals becomes significantly nonzero - and the pattern of nonzero residuals must allow isolation of the failed component. For most of this work, we have not considered the second stage of the processor which is a decision function based on some form of filtering and threshold testing to indicate the presence of the large residual. In most cases it will be clear simply from observation of the residuals whether or not the change in magnitude from nominal to failed conditions will admit a reliable test.

Generalized parity relations are a generalization of the notion of parity relations first introduced in the context of monitoring the instruments in a strapdown inertial navigation system. If, for example, one had 4 gyros measuring different components of the 3 dimensional angular velocity vector, these measurements would be linearly dependent and one could form a linear combination of them that would add to zero if all gyros were working properly. Upon failure of any one of the gyros, the residual would become distinctly nonzero. The generalization is to the construction of parity relations using possibly unlike instruments and to situations with less direct hardware redundancy. A generalized parity relation is constructed from the discrete-time model of the dynamics of a controlled system. The expressions for the measurements at a number of points in time are written, and with the use of the system transition relation the system state that appears in each expression is "backed up" to a common point in time. One can then form linear combinations of these expressions in such a way as to eliminate the state at the common time point - and thus produce an expression involving just the measurements and the control commands. This is a generalized parity relation. It uses only known data: the commands which were sent to the actuators over some window of past time, and the measurements received during that time. If the system is operating nominally, this expression adds to zero. If any of the sensors whose measurements appear in the relation should fail, or if any of the actuators whose commands appear should fail, the parity relation is no longer satisfied and the residual becomes significantly nonzero. A nonzero residual thus signals detection of a component failure. To isolate the failed component, a number of generalized parity relations must be formed - each depending on a different subset of the sensors and actuators. Enough parity relations must be monitored so that the set of large residuals uniquely identifies the failed component.

A special case of generalized parity relations are those that depend on only one sensor. These are called

"single-sensor parity relations" (SSPR). They can be constructed readily simply by using linear combinations of the expressions for the measurements from a single sensor at a number of points in time. A SSPR is therefore naturally suited to monitoring a particular sensor - although it also depends on all the actuators in general. It is not so evident that one can also construct parity relations that depend on only one actuator - plus all the sensors in general. We call these "single-actuator parity relations" (SAPR). One way to monitor all the components of a control system is to construct a SSPR for each of the sensors and a SAPR for each of the actuators. If a sensor fails, only one of the SSPRs produces a large residual while all, or most, of the SAPRs show large residuals. The decision logic then decides that one sensor has failed rather than that many actuators have failed at once. Decisions about actuators can be made in the same way because upon failure of an actuator, only one of the SAPRs shows a large residual while all, or most, of the SSPRs have large residuals.

The failure detection filter is an observer of the state of a dynamic system. It has the same structure as the Kalman filter, but its gain matrix is designed for a different purpose. Whereas the Kalman gain is designed to produce minimum variance estimates of the state variables, the detection filter gain is designed to separate the signatures of the failures of certain components into subspaces which have one dimensional projections to the output residual space. The residual space is the set of differences between the actual measurements and the filter's estimates of those measurements. The filter can be designed to monitor a set of components, and it is only the signatures of the failures of those components which are separated out. The failure of other components would typically corrupt all of the residuals just as one would expect of other observers such as a Kalman filter.

It is only the signatures of actuator failures that can be isolated to one dimension in the residual as indicated above. Sensor failures produce signatures which can be isolated to a two dimensional plane in the residual space. In the absence of any failures, the residual is small in all directions, consisting only of noise or other unmodeled effects. Upon failure of an actuator which the filter is designed to monitor, the residual grows in only one direction. That direction is indicative of which the actuator has failed. Upon failure of a sensor which the filter is designed to monitor, the residual grows in a plane which is indicative of the sensor that has failed. Upon failure of a component which the filter is not designed to monitor, the residual grows in all, or at least most, directions. This indicates failure of a component which this filter is not designed to monitor, and one should look to the residuals of another filter to isolate the failed



component.

There is an alternate way to monitor sensors which generates one dimensional signatures of the failures of all the components. That is to design two filters each using half of the available measurements. Each filter can be designed to monitor some of the actuators - whose failures will produce one dimensional signatures in the normal residuals of these filters. In addition, one can construct "cross residuals" which are the differences between one group of the measurements and the estimates of those measurements produced by the filter that does not use those measurements. So when a sensor from this group fails, that measurement is directly corrupted. But the estimate of that measurement is not corrupted because it is produced by the filter that does not use that measurement. Thus only the cross residual from that measurement becomes large. This allows one to construct a failure monitoring system in which a single residual corresponds to each of the components of the control system.

**Results** - The performance of the generalized parity relations in monitoring sensor and actuator failures on the Mini-Mast and the ACES facility will be shown first. These experiments were conducted by recording files of input and output data from the experimental facility and then simulating failures in the data files and playing that data through the failure monitor. Sensor failures can be simulated easily in this way. Several different modes of failure were simulated because the FDI processors used in this work should recognize any mode of failure. Most of the results are shown for the "sensor off" failure mode because that mode was the most difficult to see in the residual. The "sensor off" mode is the one in which the sensor puts out zero which might correspond to a number of physical problems such as loss of power, a shorted amplifier, or a broken wire. Realistic simulation of actuator failures is somewhat more difficult to do after taking the actual data. For example, if one wanted to simulate a "stuck" actuator in which the actuator does nothing in spite of being commanded to do something, the actuator cannot be simulated as doing nothing after the fact - it must be done during the taking of the data. This was done on the Mini-Mast by physically pulling a plug in the electrical circuit to interrupt the commands to the actuators for a period of time while commands were still being read into the data file. At the ACES facility this was accomplished by preparing two files of commands: the commands read into the data file had nonzero values throughout whereas the commands sent to the actuators had zeros over some interval of time.

The Mini-Mast at Langley Research Center has three actuators: the three torque wheels at the top of the mast, and many sensors. Most of these sensors are displacement

sensors which measure the displacements of different points on the mast relative to the building which houses the structure. If one were to think of these as simulating measurements which could be taken on a space structure, one would have to think of an appendage to a larger assembly with some means, possibly optical, of measuring the displacements of different points on the appendage relative to the main part of the structure. In addition to the displacement sensors, the Mini-Mast has some gyros and accelerometers.

Generalized parity relations were designed for Mini-Mast sensors and actuators in two ways: they were calculated from a given state space model of the dominant dynamics of the mast, and they were identified directly from input-output data. An example of a model-based parity relation for one of the displacement sensors is shown in Figure 24. The simulated failure mode in this case is increased variance of the sensor noise. The upper figure shows the measurement history. A noise was added to this signal from sample number 240 on. The standard deviation of this noise was 1 percent of the RMS value of the original signal; note that the noise is barely visible in the measurement. The lower figure shows the residual from the parity relation. Here the effect of this residual would threshold test on the magnitude of this residual happened in the reliably indicate that something abnormal happened in the vicinity of sample number 240. A sensitive indicator of failure is desired, of course. But this sensitivity to sensor noise also suggests difficulty in dealing with noisier sensors.

A generalized parity relation is a linear combination of some sensor outputs and actuator commands. It is a form directly suited to identification by least squares or other means. It was found that identified parity relations uniformly performed better than model-based parity relations. In addition, the model-based parity relation uses a number of lagging data points which depends directly on the order of the state space model used to construct the parity relation. When identifying a parity relation, essentially any number of lagging terms can be allowed and the identification process will determine the corresponding coefficients. Increasing the order of the parity relation was also found to be beneficial in all cases. Figure 25 shows the performance of an identified parity relation for one of the displacement sensors in the case where the sensor failed with a zero output failure mode. This parity relation also uses more terms than a model-based parity relation. The signature of the failure in the residual is clearly evident, and there would be no difficulty in making the failure decision. The corresponding model-based parity relation did not always indicate clearly a sensor-off failure. The identified parity relation of the same order





Figure 24. Top: Sensor D2 output. Noise was added to Sensor D2 from sample number 240. Bottom: Model-based SSPR  $r_2$ .




Figure 25. Identified SSPR residual  $r_3$  with 20 lags. Sensor D3 has failed to zero at sample number 234.

was more reliable, and the higher order identified parity relation, shown here, was better still.

Another degree of freedom investigated was the sampling time used in the parity relation. The baseline sampling period for the Mini-Mast system is 15 ms - and that was the sampling period used in the previous examples. Parity relations for displacement sensors perform much better with a longer sampling period. Figure 26 shows the same case as in Figure 25 except that a 30 ms sampling period was used. This means that the random sequence which drove the torque wheel actuators was held constant for 30 ms intervals and the sensor outputs were sampled at 30 ms intervals. The failure signature is much stronger in this case. So sampling period can be an important variable in the design of generalized parity relations; however, it is acknowledged that in the design of a control system the sampling period is likely to be selected on the basis of control system performance rather than FDI system performance.

Another way to improve failure signatures, especially in the case of the sensor off failure mode, is through the use of generalized parity relations based on two sensors plus all the actuators. Even when one of the involved sensors fails with zero output, these parity relations, called double sensor parity relations (DSPR), are still driven by the other sensor. This tends to produce more reliably large residuals in the failed state. Of course, a sufficient number of these DSPRs would have to be used to permit isolation of the failed sensor from the pattern of large and small residuals. Just as with SSPRs, identified parity relations performed better than model-based parity relations, and increasing the number of lags in the relation improves performance, as does increasing the sampling period.

All of the discussion so far has dealt with parity relations for the displacement sensors on the Mini-Mast. Experiments were also conducted to indicate failures of the gyro and accelerometers. The instruments involved are all located at the top of the mast. There are two accelerometers which measure lateral (bending) motion and one gyro which indicates torsional motion. All of these motions are coupled - especially in the higher frequency modes. These instruments are more sensitive to high frequency motion than are the displacement sensors, so it was considered advisable to excite the mast with signals more representative of actuator commands under closed loop operation. A simple way to reduce the bandwidth of the torque wheel commands was to use random sequences with each value held constant for 4 sampling periods. It was still clear that the parity relation residuals contained much high frequency content, so a 10 Hz. low pass filter was applied




Figure 26. Identified SSPR  $r_3$ , 20 lags, sampling period  $2T_s$ .

to the residuals before displaying them.

The filtered residuals from the y-axis accelerometer SSPR and from the z-axis gyro SSPR are shown in Figures 27 and 28. The signatures of the failures are evident, but the relative sizes of the residuals in the failed and normal states are not as large as one would like for reliable failure detection. Doubling the sampling period from 15 ms to 30 ms improves these results, but an even greater advantage results from use of double sensor parity relations in the case of these instruments. The residuals from DSPRs involving the x-axis accelerometer are shown in Figure 29 and those involving the z-axis gyro are given in Figure 30. These failure signatures are very strong. Notice that parity relations can be constructed using sensors of different types. The lower plot in Figure 29 and both plots in Figure 30 show residuals from parity relations constructed from both gyro and accelerometer data.

So clear indications of the failures of all the sensor types on the Mini-Mast were demonstrated. This was not the case with the actuators. It has been our experience generally that actuator failures are more difficult to monitor than sensor failures. To test the construction of the generalized parity relation for an actuator, a single actuator parity relation (SAPR) was constructed from a state space model of the Mini-Mast, and then a file of input and output data was generated from a simulation using the same model. An actuator off failure was then simulated in the input and output file sent to the SAPR. The resulting residual is shown in Figure 31. This residual behaves exactly as it should - being essentially zero under nominal conditions and nonzero when the actuator has failed. This test was then repeated with random noise added to the simulated measurements. The standard deviation of the noise was only 1 percent of the rms value of the measurement itself, but the effect of the noise on the parity relation residual completely obscured the failure signature - as seen in Figure 32. So it is clear that these residuals are very sensitive to sensor noise.

Another view of why actuator failures are difficult to monitor is gained by comparing the sizes of the contributions of the sensor signals and the actuator signals to the parity relations. These two components of the SAPR for the y-axis torque wheel are plotted separately in Figure 33. The actuator component, on the bottom, is far smaller than the noisy sensor component, on the top, so any abnormality in the behavior of the actuator is masked by the much larger signal coming from the sensors.

FDI experiments were also conducted using data taken on the ACES facility at the Marshall Space Flight Center at Huntsville, Alabama. A schematic diagram of that facility




Figure 27. SSPR for Y-axis accelerometer failure, 20 lags,  $1T_s$ . This residual was filtered with the postfilter. The Y-axis accelerometer has failed to zero at sample number 245.




Figure 28. SSPR for Z-axis gyro failure, 20 lags,  $1T_s$ . The residual was filtered with the postfilter. The gyro has failed to zero at sample number 255.







Figure 29. DSPR for X-axis accelerometer failure, (11,10) lags,  $1T_s$ .  
 Top: residual with sensor pair (X-axis accelerometer, Y-axis accelerometer);  
 Bottom: residual with sensor pair (X-axis accelerometer, Z-axis gyro).  
 The X-axis accelerometer has failed to zero at sample number 236.





Figure 30. DSPR for  $Z$ -axis gyro failure, (11,10) lags,  $1T_s$ . The gyro has failed to zero at sample number 286. The residual at the top was constructed from the pair ( $X$ -axis accelerometer,  $Z$ -axis gyro) and the residual at the bottom from the pair ( $Y$ -axis accelerometer,  $Z$ -axis gyro).




Figure 31. SAPR residual for Y-torque wheel failure. In this simulation the torque wheel was in a failed state between samples 213 and 284.




Figure 32. SAPR residual for Y-torque wheel failure with noisy measurements. The standard deviation of the added noise is 1% of the standard deviation of the measurement. The torque wheel was in a failed state between samples 213 and 284.





Figure 33. Contribution  $r_y$  and  $r_u$  to the SAPR residual for Y-torque wheel. The standard deviation of the noise is 1% of the standard deviation of the measurement. The torque wheel was in a failed state between samples 213 and 284.

is shown in Figure 34. It is an assembly that mimics a microwave reflector mounted on a flexible mast. It is suspended from a three-axis gimbal system at the top. Sensors include rate gyros and accelerometers at the base (top) and tip (bottom) as well as accelerometers mounted with LMED actuators at two stations along the mast. Other actuators, in addition to the LMEDs, include torque motors on the three axes of the gimbal system at the base. The experiments were conducted in essentially the same manner as those at the Mini-Mast facility.

The conclusions drawn from these experiments were exactly the same as those found in the Mini-Mast experiments. The visibility of the failure signatures was enhanced considerably by residual processing as indicated in Figure 35. The parity relation residual is initially filtered with a band pass filter to highlight the frequency range of the first few modes of the structure. These are the modes which are most excited by the applied inputs and best represented by the dynamic model inherent in a parity relation. Since we are only interested in the magnitude of this signal, it is then squared. The result is then passed through a very low frequency low pass filter, implemented as a moving average filter, which serves to stabilize the presence of the failure signature. Also, the baseline sampling rate for the ACES facility, 50 Hz., is too fast for best results with generalized parity relations. So in all the results to be shown, the inputs were held constant for 4 T and the parity relations were designed on the basis of 80 mS sampling.

No reliable state space model was available for ACES, so all the parity relations were identified directly from input-output data. The resulting performance was quite good for sensor failure monitoring. Figures 36 and 37 show the processed residuals from single-sensor parity relations for one of the base rate gyros and one of the LMED accelerometers respectively. The indications of the failures are obvious and could be reliably signaled by threshold tests on the processed residuals. As expected, double-sensor parity relations gave even stronger signatures as shown in Figures 38 and 39. These parity relations were formulated from one base gyro and the corresponding tip gyro. Double-sensor parity relations were constructed from several combinations of sensors including a base gyro and tip accelerometer, and a tip gyro and LMED accelerometer. Every combination of sensors gave reliable indications of failures.

Just as in the case of Mini-Mast, ACES actuator failures could not be monitored using generalized parity relations. The reason is the same: the contributions of the actuator command signals to the parity relations are so much smaller than the contributions of the measurements that




Figure 34. Schematic diagram of the ACES mast.




Figure 35. Block diagram of post processing filter.




Figure 36. SSPR residual for base gyro. The failure was introduced at  $t = 60$  seconds.




Figure 37. SSPR residual for LMED accelerometer. The failure was introduced at  $t = 48$  seconds.






Figure 38. DSPR residual for (base gyro, tip gyro) pair. The tip gyro failure was introduced at  $t = 50$  seconds.




Figure 39. DSPR residual for (base gyro, tip gyro) pair. The base gyro failed at  $t = 40$  seconds.

anomalies in actuator performance are completely masked by the small amount of noise in the sensor signals.

The other method of residual generation for failure monitoring which does not depend on failure mode hypotheses and is applicable to both sensors and actuators is the failure detection filter. The properties of this filter were described above. The failure detection filter requires a state space model of the system dynamics, and since a reasonably good model of the Mini-Mast was available, the experiments were conducted only with Mini-Mast data. The available model was further improved by identification; the model order used was 12. In all these experiments, the baseline sampling period of 15 ms was used. We note that in a real application, the sampling period of a control system would be chosen to achieve the required control system performance and would not be altered to improve FDI performance as was needed using generalized parity relations. So it is an important advantage of the failure detection filter that it performs well at the faster sampling rate.

It was mentioned in the description of the failure detection filter that it is most naturally suited to monitoring actuator failures in that it can restrict the signatures of actuator failures to one dimension in the residual space. This is especially fortunate in that we found that actuator failures could not be monitored using generalized parity relations. Figure 40 shows the processed residual corresponding to the X torque wheel on the Mini-Mast. The residual processing scheme shown in Figure 35 is used here. The actuator off failure mode was simulated between 30 and 60 seconds by physically disconnecting the command signal from the input to the torque wheel. The indication of that failure is obvious and would be signaled reliably by a simple threshold test on the processed residual. The corresponding processed residual for the Y torque wheel is shown in Figure 41. It gives an equally clear indication of the actuator failure. The Z torque wheel actuator gave an even stronger failure signature because the response of the Mini-Mast in torsion is greater than in bending.

It was mentioned in the description of the failure detection filter that with direct modeling of sensor failures, it is only possible to restrict their signatures to a two dimensional plane in the residual space. This is less convenient for residual monitoring than a one dimensional residual, and it is more costly in the use of the available dimensions in the residual space. Fortunately, a modification of the direct representation of sensor failures permits their failure signatures to be restricted to one dimension in the output space just as for actuators. The modification is use of state space




Figure 40. FDF residual  $r_{p1}$  for X-torque wheel failure. The residual was processed as shown in Figure 6.2




Figure 41. FDF residual  $r_{p2}$  for Y-torque wheel failure. The failure existed approximately between  $t = 30$  and  $t = 60$  seconds.

augmentation. If the sensor failure signal is supposed to be generated by some dynamic operator, and the state space model of that operator is adjoined to the state space model of the system dynamics, then the failure signal appears as a state variable rather than the direct output of the sensor. An anomaly in a state variable can be restricted to a single line in the output space of a failure detection filter. Using just first order dynamics for this purpose, the residual corresponding to the displacement sensor D1 is shown in Figure 42. This is the unprocessed residual produced directly by the failure detection filter. The failure signature is even more reliable than that produced by a single sensor generalized parity relation in that the large residual is not just a transient - it persists as long as the failure does.

Using this method to generate one dimensional signatures of sensor failures, one has the option of choosing dynamic models for the failure signals which mimic the expected behavior of the failure signals. In this case, the dominant response of the Mini-Mast to the random excitation by the torque wheels is in torsion, and the first torsion mode frequency is prominent in the measurements. For the sensor off failure mode, the failure signal is just the negative of the signal that sensor should be generating. So an oscillation at the first torsion frequency is a likely candidate for a failure signal. Use of a second order model of the failure signal generator, with natural frequency chosen near that of the first torsion mode, resulted in the unprocessed residual shown in Figure 43. The residual following the failure is not only larger in magnitude but shows clearly the expected frequency. The residual processing scheme shown in Figure 35, with the bandpass filter set to this frequency, gives a very reliable indication of the sensor failure.

So in summary, the failure detection filter performed very well in monitoring both sensors and actuators of the Mini-Mast for failures. Its indication of sensor failures was even stronger than that of single sensor generalized parity relations and its ability to monitor actuators for failures stands in contrast to the fact that generalized parity relations could not be used for that purpose at all. In addition, the failure detection filter performs well using the baseline sampling rate, which is advantageous for the control system, whereas use of longer sampling periods was important to the performance of generalized parity relations.

## Conclusions

The work under this CSI Program Guest Investigator grant centered on two aspects of the analysis, design and




Figure 42. FDF residual  $r_1$  for Sensor D1 failure.




Figure 43. FDF residual  $r_1$  for Sensor D1 failure. A second order model was used for the failure modes.

operation of control systems for large, flexible spacecraft. They are:

- The modeling of deployed or erected structures including nonlinear joint characteristics
- The detection and isolation of failures of the components of control systems for large space structures

In the first research task it was determined that use of continuum beam elements as approximations to sections of long trusses is a useful strategy to reduce the number of degrees of freedom in the representation of the dynamics of complex structures. It was also determined that sinusoidal-input describing function methodology is a useful approach to the approximate modeling of isolated nonlinear elements in structures such as joints. Taken together, these two techniques give the structural analyst a means of representing flexible structures with nonlinear properties in a way that permits analytic solutions for forced response to sinusoidal inputs and the determination of possible limit cycles in controlled structures.

The nonlinear equivalent beam finite element model was used to investigate the behavior of one and five bay 2-D trusses with various types of nonlinear joints. These include a cubic spring, both hardening and softening gain change, sliding pin, a sliding gain change, and a joint that was designed to give a structural response similar to that observed in the Mini-Mast. These results show the classic characteristics of nonlinear response such as change in resonant frequency and damping with change in amplitude of excitation. The results demonstrated the efficacy of the procedure. The methodology was also used to model the change in the behavior of a structure with nonlinear joints under gravitational loading, which puts a preload on the joints, compared with weightless conditions. This is a very important aspect of relating ground-based test results to behavior to be expected of a structure in orbit.

The Mini-Mast was also modeled using the techniques developed in this work. The model matched the frequencies of the lower frequency beam-like modes very accurately. With the nonlinear joint characteristic synthesized for this purpose, the model also matched qualitatively, at least, the nonlinear characteristics of the Mini-Mast obtained in sine sweep responses.

This modeling methodology was then used in the context of design of a control system for a structure. The amplitude-dependent variations in the properties of the structure, as determined from the model, were used to assign bounds on parameter variations for the structure. These

bounds were then used in a robust control design procedure. The resulting controller, when applied to the nonlinear structure, had significantly better disturbance rejection properties than those of the baseline controller.

In the second research task attention was centered on two methodologies for FDI residual generation which are applicable both to sensor and actuator failure monitoring, and which do not depend on assumptions about the mode of failure. These two methods are generalized parity relations and the failure detection filter. Experiments with input-output data collected on both the Mini-Mast and the ACES facility confirmed our earlier experience that it is easier to detect and isolate failures of sensors than for actuators. Both generalized parity relations and the failure detection filter could reliably monitor the performance of the sensors on both experimental facilities. In the case of Mini-Mast, this included the displacement sensors at the top of the mast and the accelerometers and z-axis gyro also mounted at the tip. In the case of ACES, this included the gyros at the base and tip, the accelerometers at the tip (bottom of the facility) and the LMED accelerometers.

Single-sensor parity relations are especially convenient for failure isolation, but it was found that single-sensor parity relations did not always give reliable failure signatures under baseline conditions. This performance was improved by use of a longer sampling period than is standard for each facility and which would be preferred for use in a control system. The failure signatures were also more reliable using double-sensor parity relations and the combination of longer sampling period and use of double-sensor parity relations gave even better performance. In every case it was true that generalized parity relations identified directly from input-output data performed better than parity relations computed from a given state space model of the system.

Generalized parity relations failed to monitor reliably the actuators in either experimental facility. By looking at the individual contributions of the sensors and the actuators to parity relations designed to isolate actuator failures it was found that the actuator contributions are much smaller than the sensor contributions. So even a relatively small amount of noise on the measurements is enough to mask anomalies in the small actuator signals and thus mask the failure signatures.

The failure detection filter requires a state space model of the system dynamics, and since no reliable model was available for ACES, it was tested on Mini-Mast data only. In that situation it performed very well. The failure detection filter could reliably monitor both the

sensors and actuators for failures. And it could work directly with the baseline sampling rate which might be needed for the purpose of a control system. So, at the cost of a somewhat greater computational burden, the failure detection filter offered performance superior to that of generalized parity relations.